How (Not) to Model Autonomous Behaviour

Ezequiel A. Di Paolo, Hiroyuki Iizuka

CSRP 588

April 2007

ISSN 1350-3162

University of Sussex

Cognitive Science
Research Papers

How (not) to model autonomous behaviour

Ezequiel A. Di Paolo[1] & Hiroyuki Iizuka[1,2]

1. Centre for Computational Neuroscience and Robotics (CCNR)
Centre for Research in Cognitive Science (COGS)
University of Sussex
Brighton, BN1 9QH, UK
ezequiel@sussex.ac.uk

2. Department of Media Architecture, Future University-Hakodate
116-2 Kamedanakano-cho, Hakodate, Hokkaido, 041-8655, Japan
ezca@sacral.c.u-tokyo.ac.jp

To appear in   , Special issue on Modelling Autonomy.

Abstract.

Autonomous systems are the result of self-sustaining processes of constitution of an identity under precarious circumstances. They may transit through different modes of dynamical engagement with their environment, from committed ongoing coping to open susceptibility to external demands. This paper discusses these two statements and presents examples of models of autonomous behaviour using methods in evolutionary robotics. A model of an agent capable of issuing self-instructions demonstrates the fragility of modelling autonomy as a function rather than as a property of a system's organization. An alternative model of behavioural preference based on homeostatic adaptation avoids this problem by establishing a mutual constraining between lower-level processes (neural dynamics and sensorimotor interaction) and higher-level metadynamics (experience-dependent, homeostatic triggering of local plasticity and re-organization). The results of these models are lessons about how strong autonomy should be approached: neither as a function, nor as a matter of external vs. internal determination.

Keywords: biological autonomy, modelling autonomous behaviour, evolutionary robotics, self-setting of goals, behavioural preference.

1. Introduction

In this paper I would like to establish two important points about autonomy that stem from a careful analysis of the continuity between life and cognition, and a third point by implication. The two main messages I would like to establish about autonomy are: 1) autonomous systems always originate in self-sustaining processes of constitution of an identity under precarious circumstances and 2) such processes can be dynamically manifested in different modes of engaging with the world ranging from committed coping to open susceptibility. The implication of these two points will be that 3) current work in "autonomous" robotics based on ideas of automated synthesis of design (e.g., evolutionary robotics) and dynamical systems approaches to cognition, is still far from achieving or even modelling autonomy in the strong sense advocated here, but that this work may be at the same time the surest route to this goal. I will concentrate for the most part of the paper on discussing examples of recent work in evolutionary robotics. One case illustrates the insufficiency of thinking about autonomy in terms of functions and another example shows that at least some interesting aspects of the organization of autonomous behaviour can be modelled fruitfully once we take points 1 and 2 more seriously. Both cases, however, constitute “good” examples of the role of modelling in clarifying complex concepts such as autonomy.

2. Why should autonomous systems generate their own identity?

I will work under the assumption that autonomous systems, i.e., systems capable in some non-trivial sense of setting their own laws, exist, and that living systems provide the clearest, less controversial examples of such autonomy (even if it may still be possible to discuss autonomous systems that are non-living; or let's say, remain agnostic about the possibility). That autonomy is not an illusion is far from evident for Western thought. This is in fact because it is often suspected to be a purely ascriptional property – one that will simply vanish upon closer inspection. Autonomy remains such a slippery concept if examined under the magnifying glass of reductionist physicalism. If we are to avoid mysteries, an autonomous system must follow only the laws of physics, hence it cannot set its own laws, therefore they don't really exist, they're just convenient ways of talking. For Kant, in his argument just too quickly takes sides in the conflict between two kinds of very real experiences: the experience of the physical world as regular and describable in terms of laws and the experience of our perceived teleology and autonomous behaviour in others and, most importantly, in ourselves. On what basis are two reliable and repeatable experiences to be discriminated as real or unreal? History tells us that this is a naive formulation and that conflict breeds novel understanding by dialectical synthesis rather than by decreeing a winner position. This is Hans Jonas's rebuttal of the Kantian lukewarm recognition of the importance, but not quite properly ontological status, of intrinsic teleology. We can know life because we ourselves are alive (Jonas, 1966; Weber & Varela, 2002; Di Paolo, 2005; Di Paolo, Rohde & De Jaegher, forthcoming).

Let's just boldly state that living organisms are autonomous – they follow laws set up by their own activity. Fundamentally, they can only be autonomous by virtue of their self-generated identity as distinct entities. A system whose identity is fully specified by a designer and cannot, by means of its own actions, regenerate its own constitution, can only follow the laws contained in its design, no matter how plastic, adaptive, or life-like its performance. In order for a system to generate its own laws it must be able to build itself *at some level of identity*. If a system 'has no say' in defining its own organization, then it is condemned to follow an externally given design like a laid down railtrack. It may be endowed with ways of changing its behaviour depending on history, but at some level it will encounter an externally imposed functional (as opposed to physical) limitation to the extent to which it can change. This can only be avoided if the

Let’s provide a definition of an autonomous system.

An

of either the environment or internal sub-agential modules meant to represent theoretical constructs such as instincts or drives. And it is only made more radical by the connection between the constitutive and interactional aspects of autonomy that is the basis of the idea of sense-making (Varela 1997; Thompson, 2007; Di Paolo, 2005), the bringing forth of a world of significance.

3. A fable about the dynamics of everyday life

When trying to understand autonomous it may be instructive to take a look at the ongoing cycles of activity in normal everyday life and how they are often very different from the performances that are studied in psychology, neuroscience, cognitive science and AI/robotics. Theur pathological action (obsessive repetition, moths attracted to the candle flame, etc.).

What happens after self-extinguished coping? It is simply contrary to everyday experience to assume that new goals will immediately follow from the attainment or frustration of previous ones (we are of course not ignoring the possibility of hierarchical organization of tasks into sub-tasks in which case the next set of activities is generally well-defined, but this is not the only possibility). In fact, our experience tells us that there are moments of certain *openness* to the possibilities afforded by our situation (such openness can clearly be very different depending on the affective outcome of the previous coping task). While distractors were robustly ignored during

Evolutionary robotics (ER) is still proving a useful and open-ended method for exploring this increasingly less constraining role of the designer that may be required to achieve strong artificial autonomy. ER hands in the task of filling in design specifications pertaining to mechanisms, morphology, structural and functional organization to an automatic process of artificial evolution (Harvey et al. 1997, Nolfi & Floreano, 2000). Thus, instead of designing a robot that must explore the environment but should go to the green light when the battery is down, one can attempt to design a robot that more generally must keep the battery up during its explorations, or more implicitly, a robot that explores indefinitely. In principle, there may be different ways of achieving this broader goal, and artificial evolution can find many of thesi84(r.4((l)-9n93.2(y)4.2(. ER)19)y9.6(al)-1(efe6(n)9.

other half the light is far from the goal. Approaching the light and remembering latching onto the *significant interactions* with the environment that will lead to achieving the desired goal efficiently and robustly. Finding a target cannot depend on the initial position of the agent, or the initial internal state, and so these parameters must be randomized from trial to trial to ascertain a level of stability of the solutions that evolve. But this very basic element of the ER methodology may play against the design of autonomous agents, at least if we consider the different dynamical regimes of activity described in the previous section. If evolution is to produce stable and robust dynamical controllers, it will avoid being strongly influenced by irrelevant environmental factors, but at the same time it will avoid internal sources of instability. Hence it will produce robust coping, but not necessarily dynamical states of openness after coping activity is self-extinguished. That's why goal-seeking evolved robots tend to keep around their targets like moths attracted to a flame. They behaviour is almost pathological. The lack of self-extinction of behaviour should perhaps be taken as a sign of bad design (cf., work by Ian Macinnes on functional circles and practical ways of dealing with this problem, e.g., Macinnes and Di Paolo, 2006). So evolving autonomous robots will have to overcome this problem by either selecting the right building blocks, or including sensorimotor interactions and internal elements that inevitably will sometimes lead to transitions between low and high dimensionality in the dynamical flow as suggested in the previous section.

5. A "self-instructing" agent. How not to model autonomy.

Let us consider an example of an agent capable of generating its own instructions and following them. In some loose sense of autonomy (but not necessarily in the operational sense that we have offered above), this agent would be setting up its own goals. I present the following agent as a computer-enhanced thought experiment but also as a demonstration of why certain tempting methodologies for designing autonomous agents are conceptually flawed. In the next section, I will show an agent that is not yet fully autonomous but which demonstrates what I consider a better methodology. Both these models demonstrate how we can learn about autonomy without yet producing proper instantiations.

In his well-known discrimination experiments, Randall Beer (2003) has shown how minimally cognitive behaviour can be 1) easily modelled and analysed using a combined evolutionary robotics/dynamical systems approach, and 2) how such models, albeit minimal, demonstrate interesting general principles and provide extendable vocabularies to discuss cognition in dynamical terms. The basic discrimination experiment consists of a visually-guided agent moving in 1 dimension (left-right) whose task is to catch a falling object if it is a circle and avoid it if it is a diamond; the agent receives input from an array of linear visual sensors (rays that activate when intersected by the falling object) and this input is fed into a recurrent, symmetrical CTRNN controller[i]. The output of the network determines the velocity of the agent, (Beer, 2003). Although dynamical analysis has shown that agents use the absolute radius of the falling shape to perform their discrimination, extensions of the setup to shapes of variable size results in agents capable of discrimination based on shape, (Di Paolo and Harvey, 2003).

Let us consider a variant of this model. An agent that performs a circle/diamond shape discrimination, but that depending on an external binary signal its choice of which object to catch can be altered. So if the external signal (ES) is set to 0, the agent is a circle-catcher and if the signal is 1, the agent is a diamond-catcher.

The setup is otherwise similar to Beer's experiments, with the difference that sensors are binary (to increase sensory ambiguity and encourage more active solutions). And additionally, a focus control is added to the array of sensor rays. This is an effector neuron that simply opens and closes the angle of the sensors rays in a linear way. Interestingly, this extra level of sensory control is important to evolve agents capable of changing their behaviour depending on the external instruction. Figure 2 shows the average fitness of 10 independent runs with and without focus control. The best focus controlling agents can perform either circle-catching or diamond-catching on demand for a relatively large range of sizes, using ambiguous noisy sensors with success rates of over 85%.

[Figure 2 about here]

These agents have now a well-defined signal that alters the goal they pursuit. Couldn't such a signal be somehow provided internally? Ideally, could such a signal be generated in a way that is jointly dependent on internal and environmental factors? Exclusive dependence on either class of factors would not generate an agent that we would be happy to call autonomous as we could suspect that the agent is following the instructions that either are external to it or is blindly taking no account of its situation. Autonomy, even in an intuitive sense, is ruled out by either of these two conditions. Why? Because both conditions negate the idea of self-determination. The case of constant reactive response to the environment is clear. No system that is simply driven externally can ever be autonomous. But, and this is less intuitive, the same may be said about a system that is "driven internally". If a subset of a system exerts control on the whole, then the situation remains that of a system that is *controlled*, not *self-determined*. If a system is controlled only by internal dynamics making it blind to the current environmental situation (what sometimes in mathematical terms is indeed called an "autonomous" system due to the lack of parametrical and time-dependent driving), the system has nothing to determine itself *against.* It simply endures in its dynamics because it's closed to environmental challenges[ii].

capable of generating a stable on-off signal depending on internal state and environmental circumstances. There are many options. One would be a central pattern generator (CPG) that oscillates with a certain frequency in the absence of input currents and settles into either a high or a low value stable attractor in the presence of input. Such a circuit can easily be hand-designed using a fully connected 2-node CTRNN (Beer, 1995) and is shown in figure 3. The CPG receives input from the visual sensors. Depending on the phase value of the oscillation orbit, the presence of input will drive the CPG to one of two possible stable fixed points (new intersections of nullclines[iii]). For one of the nodes the two fixed points correspond to high and to low firing rates respectively. This node is then connected to ES in the pre-evolved discriminator network. The agent will now produce behaviours such as those shown in figure 4. Upon repeated presentation of a circle the agent will sometimes approach it and other times avoidw[(t.4(g o)5op3.9(l)i)-.8(dwp)-3.6(u(s)12i)11-.9(g o)5op(b(ag o)5op1to)5ul9y-.000-.4(g)

implemented for each node in the CTRNN neurocontroller as 2 bands within the range of firing rates (figure 5): a low-firing and a high-firing homeostatic region (to reduce bias, the type of each region, A or B, is assigned randomly at the beginning of the evolutionary run).

[Figure 5 about here]

The idea is that if the system holds two separate (fixed) high-dimensional boxes in the space of neural dynamics which are associated with performing different behaviours, a preference could be formed by the dynamical transitions that select which box the trajectories go into and stay in. This provides a first requirement for talking about preference, that of *durability* (bottom-up construction of stability). Once a behaviour is formed, due to the stability in a box, the system keeps doing the behaviour while ignoring other behavioural possibilities. It is like a spontaneous top-down constraint that regulates the sensorimotor flow. However, some disturbances might eventually cause a breakdown of the stability and then another behaviour can be reconstructed though the homeostatic adaptive mechanisms. Since by design, the system has another region of high stability, it will be possible in the right circumstances to switch into it and then start enacting the other behavioural option. In this way, behaviour can switch due to the corresponding transitions between two boxes. One can expect to see both spontaneous and externally induced transitions from the viewpoint of the top-down and bottom-up construction or destruction of durable but impermanent dynamical modes. Here we find a second requirement, that of the possibility of *transformation*, or change in preference.

The evolved agents show interesting behaviour when two lights (A and B) are presented simultaneously in a random position. They always “chose” to go to one of the two lights, they never stay in the middle or move away from them.

while approaching light A, the lights are swapped in position to see whether the agent changes its behaviour. The result depends on the time of the swapping. If the agent is far enough, it alters its trajectory after the swap and moves towards the new position for light A. If the swap is made later, when the agent is close to light B, the agent switches to finish its approach to light B, as if its presence was now too strong a stimulus to ignore. This and similar tests indicate that a preference is maintained or changed as a combined effect of environmental factors and endogenous dynamics.

In an attempt to measure the development of a preference, agents are tested at different times during the sequence of presentations shown in figure 6 in order to find out if their choice would have been the same at that point in time if the position of the lights had been different. The distinction between a spontaneous or externally driven “decision” is made operational by observing the agent’s behaviours in different situations departing from a same initial state. If the agent “decided” to go to one of the lights endogenously, its behaviour must be robust without depending too much on environmental factors. On the contrary, if the selection were externally driven it would be affected by changes to environmental factors such as light positions (as if the agent were not “committed” enough).

Figure 7 shows the results. Each plot indicates in shades of grey the final develops towards a more stable (or committed) dynamics. It is this *alternation* between these modes – resembling the alternation between the modes of *coping* and history of interactions would be a much better way of modelling autonomy. In addition, there is only a contingent link between internal "requirements" and external interaction. Light is made relevant to the agent by a selective pressure and it is linked to an internal condition to be satisfied (homeostasis) also by evolutionary history. Organisms present much tighter double causal links between internal needs (e.g., metabolism) and sensorimotor interactions (e.g., foraging). This is again something that should be improved for a closer approach to behavioural autonomy, (Di Paolo, 2003).

7. Conclusions

What do we learn from these models? The two examples of modelling aspects of autonomous behaviour presented abvih2wn1(ral)2(l7i)-92( )]03).000.22 TD-30002 Tc.432 Tww7i)-97(nt

Ashby, W. R., 1960. Design for a brain: The origin of adaptive behaviour (Second edition). London: Chapman and Hall.

Beer, R. D., 1995. On the dynamics of small continuous-time recurrent neural networks. Adaptive Behavior, 3, 471 – 511.

Beer, R. D., 2003. The dynamics of active categorical perception in an evolved model agent. Adaptive Behavior, 11, 209–243.

Brooks, R. A., 1991. Intelligence without representation. Artificial Intelligence, 47, 139– 159.

Di Paolo, E. A., 2000. Homeostatic adaptation to inversion of the visual field and other sensorimotor disruptions. In Meyer, J.-A., Berthoz, A., Floreano, D., Roitblat, H., & Wilson, S. (Eds.), From Animals to Animats 6: Proceedings of the Sixth International Conference on the Simulation of Adaptive Behavior. Cambridge MA: MIT Press.

Di Paolo, E. A., 2003. Organismically-inspired robotics: Homeostatic adaptation and natural teleology beyond the closed sensorimotor loop. In Murase, K., & Asakura, T. (Eds.), Dynamical Systems Approach to Embodiment and Sociality, pp. 19–42. Advanced Knowledge Internaltional, Adelaide.

Di Paolo, E. A., 2005. Autopoiesis, adaptivity, teleology, agency. Phenomenology and the Cognitive Sciences, 4, 429–452.

Di Paolo E. A., & Harvey, I., 2003. Decisions and noise: The scope of evolutionary synthesis and dynamical analysis. Adaptive Behavior, 11, 284 - 288.

Di Paolo, E. A., Rohde, M., & De Jaegher, H., Forthcoming. Horizons for the enactive mind: Values, social interaction and play. In O. Gapenne, J. Stewart and E. A. Di Paolo (eds) Enaction: towards a new paradigm in cognitive science. Cambridge MA: MIT Press.

Goldstein, K., 1995/1934. The organism. New York: Zone Books.

Harvey, I., Di Paolo, E. A., Wood, R., Quinn, M., & Tuci, E., 2005. Evolutionary robotics: A new scientific tool for studying cognition. Artificial Life, 11, 79–98.

Harvey, I., P., H., Cliff, D., Thompson, A., & Jakobi, N., 1997. Evolutionary robotics: the Sussex approach. Robotics and Autonomous Systems, 20, 207 – 224.

Iizuka, H. & Di Paolo E. A., Forthcoming. Towards Spinozist robotics: Exploring the minimal dynamics of behavioural preference. Adaptive Behavior.

Jonas, H., 1966. The phenomenon of life: Towards a philosophical biology. Evanston, IL: Northwestern University Press.

Kelso, J. A. S., 1995. Dynamic patterns: The self-organization of brain and behavior. Cambridge MA: MIT Press.

Macinnes, I., & Di Paolo, E. A., 2006. The advantages of evolving perceptual cues Adaptive Behavior 14(2), 147-156.

Moreno, A., & Etxeberria, A., 2005. Agency in natural and artificial systems. Artificial Life, 11, 161–176.

Nolfi, S., & Floreano, D., 2000. Evolutionary robotics. The biology, intelligence, and technology of self–organizing machines. Cambridge MA: MIT Press.

Rohde, M., & Di Paolo, E. A., 2006. 'Value signals': An exploration in evolutionary robotics. Cognitive science research paper 584, COGS, University of Sussex.

Rutkowska, J. C., 1997. What's value worth? Constraints on unsupervised behaviour acquisition. In Husbands, P., & Harvey, I. (Eds.), Proceedings of the Fourth European Conference on Artificial Life, pp. 290 – 298. Cambridge, MA: MIT Press.

Smithers, T., 1997. Autonomy in robots and other agents. Brain and Cognition, 34, 88 – 106.

Thelen, E., & Smith, L. B., 1994. A dynamic systems approach to the development of cognition and action. Cambridge, MA: MIT Press.

Thompson, E., 2007. Mind in life: Biology, phenomenology, and the sciences of mind. Cambridge, MA: Harvard University Press.

Tuci, E., Quinn, M., & Harvey, I., 2003. An evolutionary ecological approach to the study of learning behavior using a robot based model. Adaptive Behavior, 10, 201 – 222.

Varela, F. J., 1979. Principles of biological autonomy. New York: Elsevier, North Holland.

Varela, F. J., 1991. Organism: A meshwork of selfless selves. In Tauber, A. I. (Ed.), Organism and the origin of the self, pp. 79 – 107. Netherlands: Kluwer Academic.

List of figures

Figure 1. Dynamical modes describing the flow of everyday activity.

Figure 2. Fitness achieved for agents evolved to catch circles or diamonds under external instruction with and without focus control of sensor rays. Average of 10 independent runs.

Figure 3. Top left: CTRNN neurocontroller for self-instructing agent. ES: external signal, F: focus effector, ML and MR, motor neurons driving left and right respectively. Top right: nullclines corresponding to fully connected 2-node CPG in the absence of input. Bottom left and right: nullclines in the presence of input, trajectory ends in a low firing fixed point for neuron 1 (left) or in a high firing fixed point (right) depending on phase. Output of neuron 1 is fed into ES.

Figure 4. Repeated presentation of falling circles (left) and diamonds (right) for self-instructing agent. Plots show the horizontal displacement of the agent over time and the position where the objects fall. Agent sometimes approaches the target, other times avoids it.

Figure 5. Left: schematic representation of two high-dimensional homeostatic regions in the space of neural firing rates. Right: how the homeostatic regions are implemented for each node in the network. The plot indicates the plasticity function ( ) as a function of neural firing rate ( ). Changes to incoming weights are calculated as a function of pre- and post-synaptic activation multiplied by . Whenever the post-synaptic firing rate is in one of the two flat regions, = 0 and local plasD.009e476.64 347(p)-4(o7(l)-13.37(pl)-1)-4(o7(y)(o7(is.37(p)6.11(r)1.9(h.7(i)-4(o7(e)8ca

Figure 1. Dynamical modes describing the flow of everyday activity.

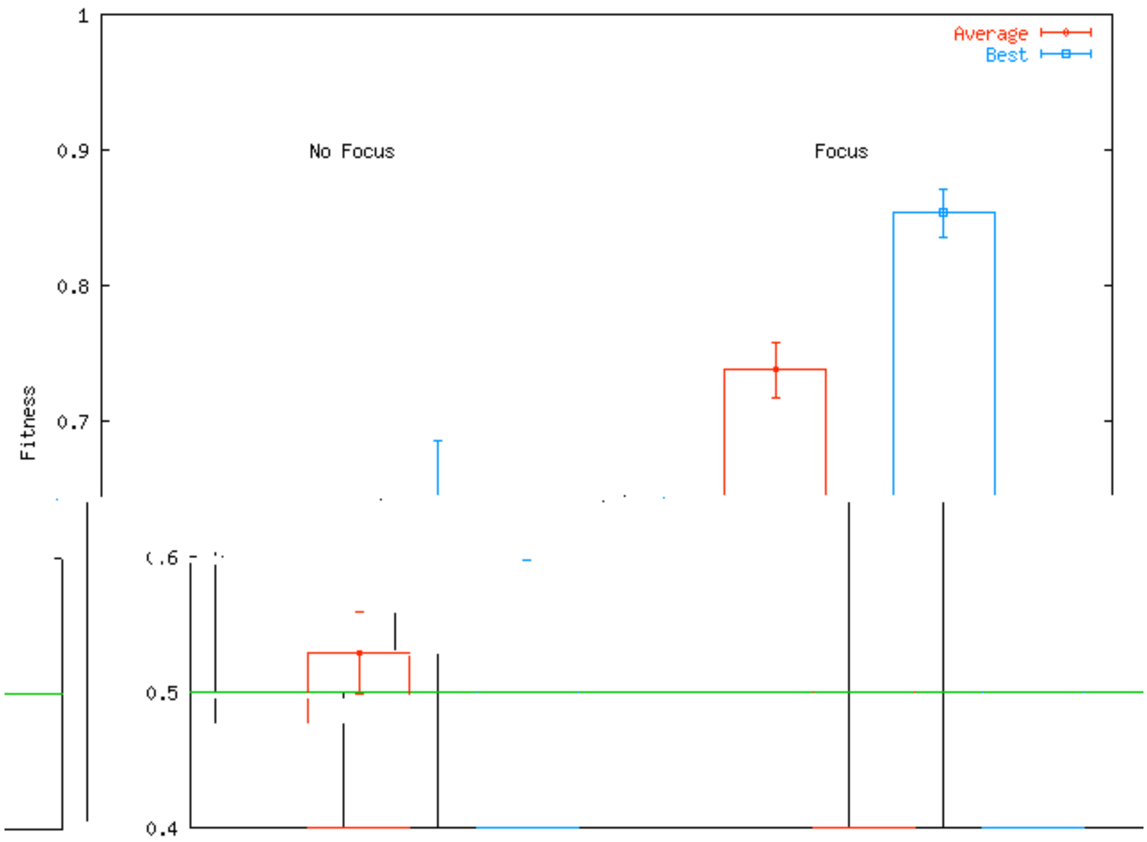
Average
Best
No Focus
Focus
Fitness
1
0.9
0.8
0.7
0.6
0.5
0.4

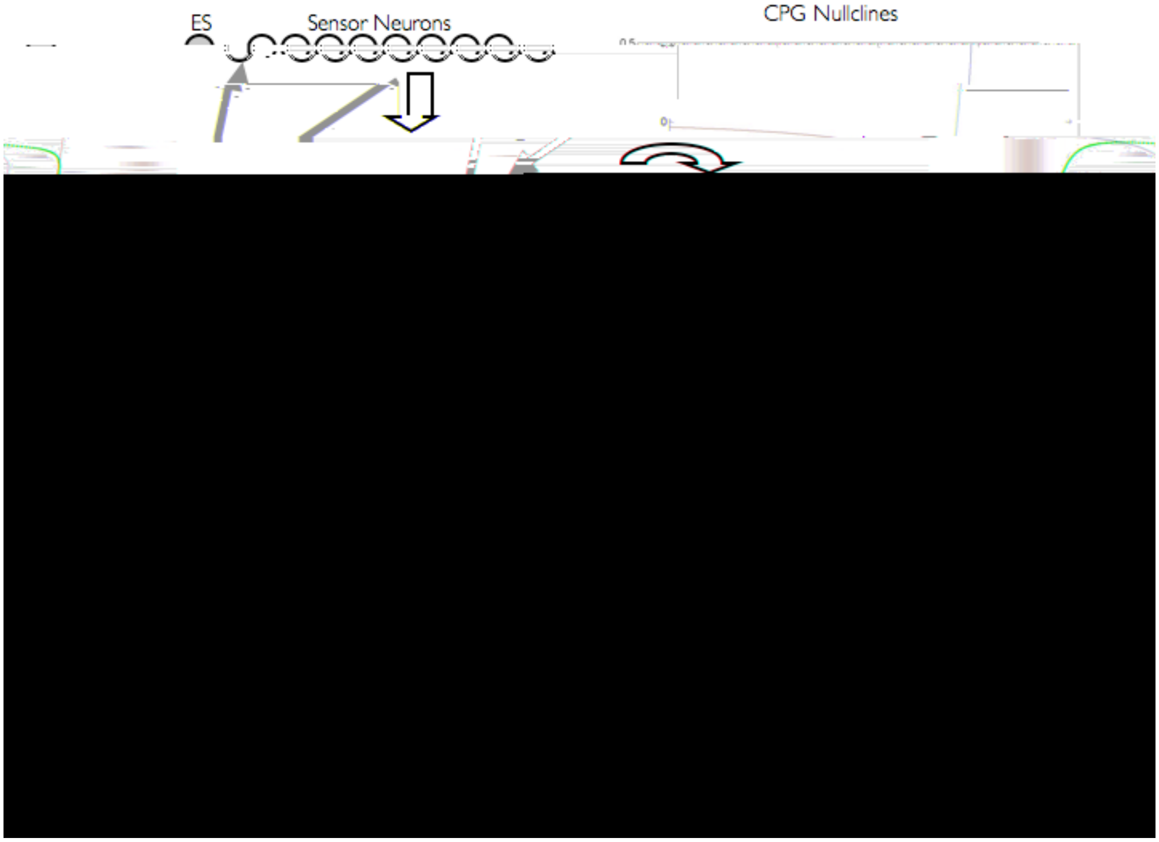

Figure 3. Top left: CTRNN neurocontroller for self-instructing agent. ES: external signal, F: focus effector, ML and MR, motor neurons driving left and right respectively. Top right: nullclines corresponding to fully connected 2-node CPG in the absence of input. Bottom left and right: nullclines in the presence of input, trajectory ends in a low firing fixed point for neuron 1 (left) or in a high firing fixed point (right) depending on phase. Output of neuron 1 is fed into ES.

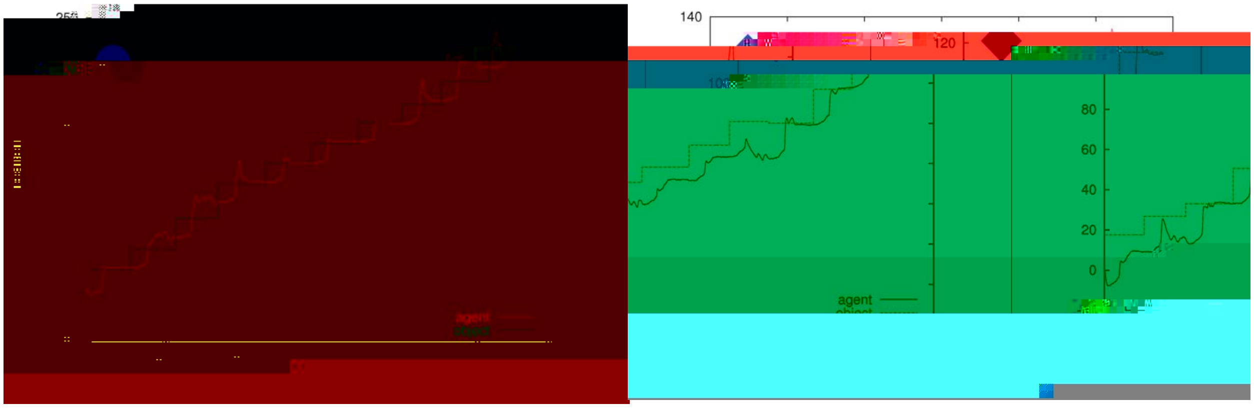

Figure 4. Repeated presentation of falling circles (left) and diamonds (right) for self-instructing agent. Plots show the horizontal displacement of the agent over time and the position where the objects fall. Agent sometimes approaches the target, other times avoids it.

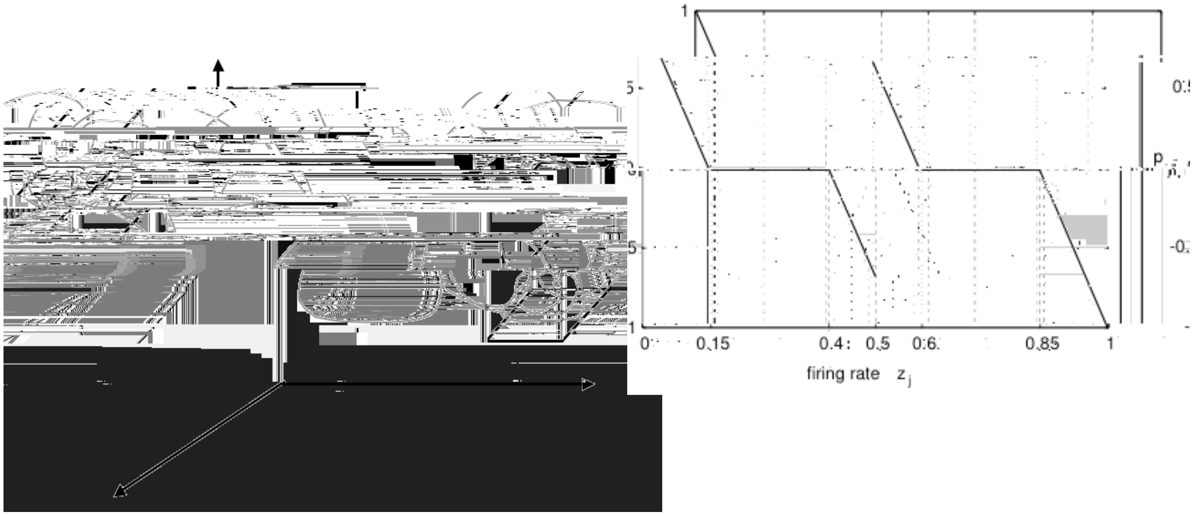

Figure 5. Left: schematic representation of two high-dimensional homeostatic regions in the space of neural firing rates. Right: how the homeostatic regions are implemented for each node in the network. The plot indicates the plasticity function (p_j) as a function of neural firing rate (z_j). Changes to incoming weights are calculated as a function of pre- and post-synaptic activation multiplied by p_j: $w_{ji} = \; _{ji}\, z_i\, p_j\, (z_j)$ where $_{ji}$ is an evolved constant. Whenever the post-synaptic firing rate is in one of the two flat regions, $p_j = 0$ and local plasticity is inhibited.

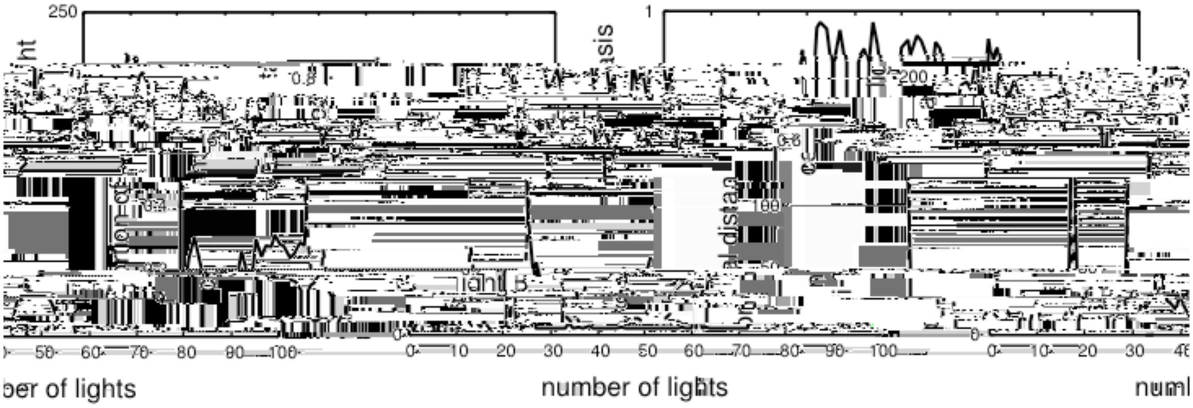

Figure 6. Left: Final distance to each light at the end of trials on serial presentations of 100 pairs of lights. Right: Proportion of neurons that have stayed within the homeostatic region for each light in correspondence to trials on the left. Adapted from Iizuka and Di Paolo (forthcoming).

Figure 7. Light preference of the agent corresponding to the states of (a) 20, (b) 25, (c) 50 or (d) 95 in Fig. 6, against different light positions. Horizontal and vertical axes indicate the initial angles of lights A and B relative to the agent's orientation respectively. The positions of lights whose difference is less than /2 are removed in order to better determine which light the agent is approaching. The dark grey circles show that the agent approaches light A. The light grey circles correspond to light B and black shows the agent does not approach either of lights. Adapted from Iizuka and Di Paolo (forthcoming).

Notes

---

[i] The state equation for a CTRNN neuron is:

$$\%_i \, (dy_i / dt) = -v_i + \,!_j \, w_{ji} z_j + I_i,$$

where i indexes all neurons, j indexes all links inputting to neuron i (which may be an empty set), $\%_i$ is a time constant, y_i is the neuron state (analogous to a membrane potential), I_i